

A PERSPECTIVE ON THE FUTURE OF THE MAGNETIC HARD DISK DRIVE (HDD) TECHNOLOGY

Rajat Chaudhary¹, Archika Kansal²

¹Asst. Professor, Department of Information Technology, Bharat Institute of Technology, Meerut

²Department of Information Technology

¹rajatlibran@gmail.com, ²archikakansal94@gmail.com

Abstract- It is estimated that over 80% of all new information produced in the world is being stored on storage devices such as pen-drive, storage chips or magnetic media, but most of it on hard disk drives. This paper presents a detailed introduction of the working, components and logical operations of storage device especially we focuses on the magnetic disk drive i.e. Hard Disk. In addition to presenting failure statistics, we analyze the correlation between failures and several parameters generally believed to impact longevity. Our analysis identifies several parameters from the drive's self monitoring facility (SMART) that correlate highly with failures. Despite this high correlation, we conclude that models based on SMART parameters alone are unlikely to be useful for predicting individual drive failures. Finally, future of information storage is examined, and storage technologies toward 1TB recording are investigated and discuss about how can we increase the performance in the hard disk drive.

Index Terms— HDD, SSD, Magneto-resistive head, Platters, Actuator, Spindle, Voice Coil, SATA, Skyrmions.

I. WHAT IS A HARD DISK DRIVE (HDD) ?

A hard disk drive (HDD) is a data storage device used for storing and retrieving digital information using rapidly rotating disks (platters) coated with magnetic material. An HDD retains its data even when powered off. Data is read in a random- access manner, meaning individual blocks of data can be stored or retrieved in any order rather than sequentially. An HDD consists of one or more rigid ("hard") rapidly rotating disks (platters) with magnetic heads arranged on a move the moving actuator arm to read and write data to the surfaces. The Hard drive uses two important principles about the magnetic fields. When we write data onto the hard drive it uses the law of electromagnetic induction and some material is magnetic and when we read data from the hard drive, it uses Lenz's law.

II. INTRODUCTION OF HARD DISK DRIVE

As today's modern world is of automation and manual storing of all the important data is day by day keeps on decreasing. So, in automation all the storage, retrieval, manipulation of the information or data is done based on peripheral devices such as laptops, computers, mobile phones etc. As hard disk has the capacity to store maximum amount of data comparative to the other storage devices like memory chips, pen drives etc. This paper depicts mainly on the internal logical working process of the storage and the components of the HDD.

The Hard Disk was invented in 24 December, 1954 developed by IBM team led and developer is Red Johnson became the dominant secondary storage device for general-purpose computers by the early 1960s. Continuously improved, HDDs have maintained this position into the

modern era of servers and personal computers. More than 200 companies have produced HDD units, though most current units are manufactured by Seagate, Toshiba and Western Digital. Figure below is an example of the Hard Disk Drive.



Fig. 1: A 2.5-inch SATA hard drive



Fig. 2: A disassembled and labelled 1997 HDD lying atop a mirror

Most people are amazed when they discover they can store hundreds of CD's worth of music on an iPod digital music player no bigger than a pack of cards. An iPod (one of the older ones, anyway) is no much more than a hard drive: an incredibly efficient computer memory device that uses simple magnetism to store vast amounts of information. The microprocessor in your computer is the bit that does all the "thinking" and calculating- but it's the hard drive that gives your computer its prodigious memory and lets you store digital photos, music files, and text documents.

III. EXISTING SYSTEM VS PROPOSED SYSTEM

Magnetic hard disk drives have significantly improved in size, performance and cost due to many technological innovations; including magneto resistive heads, low noise thin film disks, PRML channels and advanced mechanical actuators and motors. By analyzing specification trends of each new disk drive design, a perspective of this evolution can be developed and design characteristics of future disk drives can be estimated. Areal density increases have

IV. CHARACTERISTICS OF AN HARD DISK DRIVE

exceeded the traditional semiconductor development trajectory and have yielded higher-capacity, higher-performance, and smaller-form-factor disk drives, enabling desktop and mobile computers to store multi-gigabytes of data easily. Server systems containing large numbers of drives have achieved unparalleled reliability, performance, and storage capacity. All of these characteristics have been achieved at rapidly declining disk costs.

This paper relates advances in disk drives to corresponding trends in storage systems and projects where these trends may lead in the future. The progress of hard disk drive technology is reviewed, and future storage technology is observed. First, areal density growth, and improvement of performance such as data transfer rate and access speed are looked at, corresponding with the progress of magnetic materials. Then, state of the art high-density recording technologies are reviewed. The bit aspect ratio is considered with respect to the trends on precision process of pole tips, and high output sensor evolution. In the existing system the problem with the performance of the existing hard disk drive is that as it heats up too much some amount of the data loss is a big factor that affects the performance. The proposed system is to increase the capacity i.e. size and data transfer speed also data access speed of the hard disk drive and also we had discuss about a new particle to increase the performance in the future.

A groundbreaking new particle has been discovered that could potentially be used in extremely dense hard drives, known as SKYRMIONS or microscopic twisted magnetic vortices, the particle could prove extremely useful in the future of storage on magnetic hard drives, where traditional drives fast approach a density barrier.

The primary characteristics of an HDD are its capacity and performance. **Capacity** is specified in unit prefixes corresponding to powers of 1000: a 1-terabyte (TB) drive has a capacity of 1,000 gigabytes (GB; where 1 gigabyte = 1 billion bytes). Typically, some of an Hard disk drives capacity is unavailable to the user because it is used by the file system and the computer operating system, and possibly inbuilt redundancy for error correction and recovery. **Performance** is specified by the time required to heads to a track or cylinder (average access time) plus the time it takes for the desired sector to move under the head (average latency, which is a function of the physical rotational speed in revolutions per minute), and finally the speed at which the data is transmitted (data rate).

The two most common form factors for modern Hard disk drives are 3.5-inch in desktop computers and 2.5-inch in laptops. Hard disk drives are connected to systems by standard interface cables such as SATA (Serial ATA), USB or SAS (Serial attached SCSI) cables. As of 2014, the primary competing technology for secondary storage is flash memory in the form of solid-state drives (Secondary storage drives). Hard disk drives are expected to remain the dominant medium for secondary storage due to predicted continuing advantages in recording capacity, price per unit of storage, write latency and product lifetime. However, Secondary storage drives are replacing Hard disk drives where speed, power consumption and durability are more important considerations.

Improvement of HDD characteristics over time			
Parameter	Started with	Developed to	Improvement
Capacity (formatted)	3.75 megabytes	eight terabytes	two-million-to-one
Physical volume	68 <u>cubic feet</u> (1.9 <u>m³</u>)	2.1 <u>cubic inches</u> (34 <u>cc</u>)	57,000-to-one
Weight	2,000 <u>pounds</u> (910 <u>kg</u>)	2.2 <u>ounces</u> (62 <u>g</u>)	15,000-to-one
Average access time	about 600 milliseconds	a few milliseconds	About 200-to-one
Price	US\$9,200 per megabyte-	< \$0.05 per gigabyte by 2013	180-million-to-one
Areal density	2,000 <u>bits per square inch</u>	826 gigabits per square inch in 2014	> 400-million-to-one

In the above table we understand the improvement in the capacity and performance of the HDD. Early in 1956 when the Hard disk drives were introduced as data storage for an IBM real-time transaction processing computer and were developed for use with general-purpose mainframe and minicomputers. The first IBM drive, the 350 RAMAC, was approximately the size of two refrigerators and stored 5 million 6-bit characters (3.75 megabytes) on a stack of 50 disks. In 1962 IBM introduced the model 1311 disk drive, which was about the size of a washing machine and stored two million characters on a removable disk pack then later in 1980s reached capacities of 300 megabytes. Non-removable Hard disk drives were called "fixed disk" drives. Some high performance Hard disk drives were manufactured with one head per track, e.g., IBM 2305 so that no time was lost

physically moving the heads to a track. Known as Fixed-Head or Head-Per-Track disk drives they were very expensive and are no longer in production.

External Hard disk drives remained popular for much longer on the Apple Macintosh. Every Mac made between 1986 and 1998 has a SCSI port on the back, making external expansion easy; also, "toaster" Compact Macs did not have easily accessible HDD bays (or, in the case of the Mac Plus, any hard drive bay at all), so on those models, external SCSI disks were the only reasonable option. Driven by ever increasing areal density since their invention, Hard disk drives have continuously improved their characteristics. At the same time, market application expanded from mainframe computers of the late 1950s to most mass storage applications including computers and consumer applications such as storage of entertainment content.

V. HOW TO STORE INFORMATION WITH MAGNETISM

An HDD records data by magnetizing a thin film of ferromagnetic material on a disk. Sequential changes in the direction of magnetization represent binary data bits. The data is read from the disk by detecting the transitions in magnetization. User data is encoded using an encoding scheme, such as run-length limited encoding which determines how the data is represented by the magnetic transitions. A typical HDD design consists of a spindle that holds flat circular disks, also called platters, which hold the recorded data. The platters are made from a non-magnetic material, usually Aluminium alloy, glass or ceramic and are coated with a shallow layer of magnetic material typically 10–20 nm in depth, with an outer layer of carbon for protection.

If our computer has a 20 gigabyte (GB) hard drive, or you have a 20 GB iPod or MP3 player, it's a bit like a box containing 1.6 million microscopically small iron nails, each of which can store one tiny piece of information called a bit. A bit is a binary digit—either a number zero or a number one. In computers, numbers are stored not as decimal (base-10) but as patterns of binary digits instead. For example, the decimal number 382 is stored as the binary number 101111110. Letters and other characters can also be stored as binary numbers. Thus, computers store a capital letter A as the decimal number 65 or the binary number 1000001. Suppose you want to store the number 1000001 in your computer in that big box of iron nails. You need to find a row of seven unused nails. You magnetize the first one (to store a 1), leave the next five demagnetized (to store five zeros), and magnetize the last one (to store a 1).

HARD DRIVE INTERNAL ASSEMBLY PHOTO



Figure 3: Diagram labelling the major components of a HDD

Recording of single magnetization of bits on a 200 MB HDD-platter (recording made visible using CMOS-Magnetic View)




Figure 4: Longitudinal recording (standard) & perpendicular recording diagram Parts of the Hard Disk Drive

HARD DRIVE INTERNAL PARTS DIAGRAM




Figure 5: HDD Parts

A hard drive has only a few basic parts. There are one or more shiny silver platters where information is stored magnetically, there's an arm mechanism that moves a tiny magnet called a read-write head back and forth over the platters to record or store information, and there's an electronic circuit to control everything and act as a link



between the hard drive and the rest of your computer. In this diagram shows the different parts of the hard Disk Drive.

A. Components of the Hard Disk Drive

1. Actuator: that moves the read-write arm.
2. Read-write arm: swings read-write head back and forth across platter.

3. Central spindle: allows platter to rotate at high speed.
4. Magnetic platter: stores information in binary form.
5. Plug connections: link hard drive to circuit board in personal computer.
6. Read-write head: is a tiny magnet on the end of the read-write arm.

7. Circuit board: on underside controls the flow of data to and from the platter.
8. Flexible connector: carries data from circuit board to read-write head and platter.
9. Small spindle: allows read-write arm to swing across platter.

B. The Platters




Figure 6: Platters of the HDD

The platters are the actual disk in the drive that store the magnetized data. Traditionally Platters are made up of a light aluminium alloy and coated with a magnetizable material such as ferrite compound that is applied in liquid form and spun evenly across the platter or thin metal film plating that is applied to the platter through electroplating. Newer technology uses glass and/or ceramic platters because they can be made thinner and also because they are more efficient at resisting heat. The magnetic layer on the platters has tiny domains of magnetization that are oriented to store information that is transferred through the read/write head. Most drives have at two platters, and the larger the storage capacity of the drive, the more platters there are. Each platter is magnetized on each side, so a drive with 2 platters has 4 sides to store data.

The platters in contemporary Hard disk drives are spun at speeds varying from 4,200 rpm in energy-efficient portable devices, to 15,000 rpm for high-performance servers. The first Hard disk drives spun at 1,200 rpm and, for many years, 3,600 rpm was the norm. As of December 2013, the platters in most consumer-grade Hard disk drives spin at either 5,400 rpm or 7,200 rpm. Information is written to and read from a platter as it rotates past devices called read-and-write heads that operate very close (often tens of nano meters) over the magnetic surface. The read-and-write head is used to detect and modify the magnetization of the material immediately under it.

C. The spindle and spindle motor

In modern drives there is one head for each magnetic platter surface on the spindle, mounted on a common arm. An actuator arm (or access arm) moves the heads on an arc

(roughly radially) across the platters as they spin, allowing each head to access almost the entire surface of the platter as it spins. The arm is moved using a voice coil actuator or in some older designs a stepper motor. Early hard disk drives wrote data at some constant bits per second, resulting in all tracks having the same amount of data per track but modern drives (since the 1990s) use zone bit recording—increasing the write speed from inner to outer zone and thereby storing more data per track in the outer zones. In modern drives, the small size of the magnetic regions creates the danger that their magnetic state might be lost because of thermal effects, thermally induced magnetic instability which is commonly known as the "super paramagnetic limit." To counter this, the platters are coated with two parallel magnetic layers, separated by a 3-atom layer of the non-magnetic element ruthenium, and the two layers are magnetized in opposite orientation, thus reinforcing each other. Another technology used to overcome thermal effects to allow greater recording densities is perpendicular recording, the technology was used in many Hard disk drives.

The platters in a drive are separated by disc spacers and are clamped to a rotating spindle that turns all the platters in unison. The spindle motor is built right into the spindle or mounted directly below it and spins the spindle at a constant set range ranging from 3,600 to 7,200 RPM. The motor is attached to a feedback loop to ensure that it spins at precisely the speed is supposed to. HDD with disks and motor hub removed exposing copper colored stator coils surrounding a bearing in the center of the spindle motor. Orange stripe along the side of the arm is thin printed-circuit cable, spindle bearing is in the center and the actuator is in the upper left.



A typical HDD has two electric motors: a spindle motor that spins the disks and an actuator (motor) that positions the read/write head assembly across the spinning disks. The disk motor has an external rotor attached to the

disks; the stator windings are fixed in place. Opposite the actuator at the end of the head support arm is the read-write head; thin printed-circuit cables connect the read-write heads to amplifier electronics mounted at the pivot of the actuator.

The head support arm is very light, but also stiff; in modern drives, acceleration at the head reaches 550 g.

D. Actuator Coil

The actuator coil motor that moves the read-write arm. In older hard drives, the actuators were stepper motors. In most modern hard drives, voice coils are used instead. As their name suggests, these are simple electromagnets, working rather like the moving coils that make sounds in loudspeakers. They position the read-write arm more quickly, precisely, and reliably than stepper motors and are less sensitive to problems such as temperature variations.




Figure 7: Head stack with an actuator coil on the left and read/write heads on the right

The actuator is a permanent magnet and moving coil motor that swings the heads to the desired position. A metal plate supports a squat neodymium-iron-boron (NIB) high-flux magnet. Beneath this plate is the moving coil, often referred to as the voice coil by analogy to the coil in loud speakers, which is attached to the actuator hub, and beneath that is a second NIB magnet, mounted on the bottom plate of the motor (some drives only have one magnet).

The voice coil itself is shaped rather like an arrowhead, and made of doubly coated copper magnet wire. The inner layer is insulation, and the outer is thermoplastic, which bonds the coil together after it is wound on a form, making it self-supporting. The portions of the coil along the two sides of the arrowhead (which point to the actuator bearing center) interact with the magnetic field, developing a tangential force that rotates the actuator. Current flowing radially outward along one side of the arrowhead and radially inward on the other produces the tangential force. If the magnetic field were uniform, each side would generate opposing forces that would cancel each other out. Therefore the surface of the magnet is half N pole, half S pole, with the radial dividing line in the middle, causing the two sides of the coil to see opposite magnetic fields and produce forces that add instead of cancelling. Currents along the top and bottom of the coil produce radial forces that do not rotate the head.

The Hard disk drives electronics control the movement of the actuator and the rotation of the disk, and perform reads and writes on demand from the disk controller. Feedback of the drive electronics is accomplished by means of special segments of the disk dedicated to servo feedback. These are either complete concentric circles (in the case of dedicated servo technology), or segments interspersed with real data (in the case of embedded servo technology). The servo feedback optimizes the signal to noise ratio of the GMR sensors by adjusting the voice-coil of the actuated arm. The spinning of the disk also uses a servo motor. Modern disk firmware is capable of scheduling reads and writes efficiently on the platter surfaces and remapping sectors of the media which have failed.

E. Reading and writing data

HARD DRIVE DATA READ & WRITE OPERATION MOTION DIAGRAM



Figure 8: HDD operation motion

The most important thing about memory is not being able to store information but being able to find how to store it. Imagine storing a magnetized iron nail in a pile of 1.6 million identical nails and you'll have some idea how much trouble your computer would get into if it didn't use a very methodical way of filing its information. When your computer stores data on its hard drive, it doesn't just throw magnetized nails into a box, all jumbled up together. The data is stored in a very orderly pattern on each platter. Bits of data are arranged in concentric, circular paths called **tracks**. Each track is broken up into smaller areas called **sectors**. Part of the hard drive stores a map of sectors that have already been used up and others that are still free. (In Windows, this map is called the **File Allocation Table** or **FAT**.) When the computer wants to store new information, it takes a look at the map to find some free sectors. Then it instructs the read-write head to move across the platter to exactly the right location and store the data there. To read information, the same process runs in reverse.

With so much information stored in such a tiny amount of space, a hard drive is a remarkable piece of engineering. That brings benefits (such as being able to store 500 Compact drives on your iPod)—but drawbacks too. One of them is that hard drives can go wrong if they get dirt or dust inside them. A tiny piece of dust can make the read-write head bounce up and down, crashing into the platter and damaging its magnetic material. This is known as a **disk crash** (or **head crash**) and it can (though it doesn't always) cause the loss of all the information on a hard drive. A disk crash usually occurs out of the blue, without any warning. That's why you should always keep backup copies of your important documents and files, either on another hard drive, on a compact disc (CD) or DVD, or on a flash memory stick.



Figure 9: Read & Write Head of HDD

VI. THE HARD DRIVE LOGICAL OPERATION DETAILS

As shown in the Hard Drive Data Read & Write Operation Motion Diagram above, the Platters spin and the Actuator Arm moves the Heads of each of the platters to Read and Write data to the Hard Drive Platters. Data is read and written via a logical process of surface electromagnetic manipulation of the Hard Drive Platters through the operation of the Heads as controlled by the Printed Circuit Board. The Printed Circuit Board (PCB) on most hard drives is located on the outside, bottom of the hard drive assembly and also most susceptible to the outside environment of your computer's case/enclosure. 90% of most Hard Drive Failures can be directly or indirectly traced back to PCB faults and

damage, many times just requiring the exact replacement of the hard drive's PCB to regain operation. The Printed Circuit Board contains electronic integrated circuit chips which have functions contained within them in the form of logical/software instructions used to manipulate and operate the physical parts of the hard drive, control how data is read and written on the Hard Drive Platters, and interpret the data back into operating system (Windows), various programs (like word processors, photo editor programs, video players, etc.) and user data files, ultimately getting this information to and from the hard drive via an interface (logical/electrical connection) within your computer's motherboard, memory (RAM) and processor (CPU i.e. Intel or AMD chip) which displays all this on your computer's screen.



Figure 10: Single platter surface of HDD and Voice Coil

The Hard Drive Data (Single Platter) Surface Detail Diagram above, shows how the Heads of each of the Platters Read and Write data to the Hard Drive Platters, in the form of Cylinders, Tracks, Clusters and Sectors. These are the logical parts of your hard drive that dictate how data is arranged and stored on your hard drive's Platters. The Hard Drive Data (Single Platter) Surface Detail Diagram above just shows the top most hard drive Platter and how data is arranged/stored on it. The Hard Drive Platters assembly, layout of Sectors, Clusters, Tracks, Cylinders, positioning and movement of the Actuator Arm and Head, are all precision based down to 1/1000th of an inch. Any part being out of alignment by more than that 1/1000th of an inch can also cause various failures. Next to your Computer's CPU (Intel/AMD processor), the hard drive is the second most highly precision built component of your computer. The average hard drive contains four disk platters (each stacked on top of the other and spaced evenly using precision made spacers on the Hard Drive Spindle. The above diagram is more of a "2 Dimensional View" of how data is stored on a hard drive, but in actuality data is really stored in more of a "3 Dimensional View" as shown in the Hard Drive (multi Platter) Assembly Surface Data Blow up Diagram.



Figure 10: Multi Platter Surface of HDD

Any one file on your computer's hard drive, may either be written sequentially in a group of Sectors that make up a Cluster (where multiple Clusters make up a Track) in a particular Track (where multiple Tracks makeup a Cylinder), or the file's data may be scattered (due to hard drive fragmentation) over various Sectors, Clusters, Tracks and Cylinders, non-sequentially. Non-sequentially stored data on your hard drive takes longer to fully load/open as opposed to sequentially stored data. Through normal use of your computer, the hard drive will generally start to write more data non-sequentially over time, as various other data read

and writes to the hard drive may create open Sectors (gaps) in Tracks on the hard drive Platters, this leads to what is termed Fragmentation (where files/data become scattered fragments on your computer's hard drive. High Data Fragmentation can lead to Logical Errors where data becomes accidentally corrupted and/or mixed together with other data necessary for other files, causing hard drive Read/Write or I/O Failures. Additionally, corruptions in your computer's operating system (Windows), due to viruses and other faults, can lead to larger Logical Errors and Total Logical failures (inability to read any data from hard drive) and can further destroy data.

VII. THE HARD DRIVE FAILURE DIAGNOSTIC PROCESS

Based on the above details about the Physical/Mechanical operation and Logical operation of your computer's hard drive, you can guess that the Hard Drive Failure Diagnostic Process is broken down by Physical/Mechanical Failure and Logical Faults/Failures with each type of failure having further subsets of faults and failures based on the lowest levels/components involved. Typical Hard Drive Failure Symptoms described by the owners include:

- a) Noise: Clicking, grinding, squealing.
- b) Heat: Hot hard drive enclosure top, bottom or sides.
- c) Smoke: Smoke coming from PCB or inside of drive.
- d) Access: Can't see hard drive listed in My Computer.
- e) Operation: Disk works intermittently or stops/starts.
- f) Errors: Errors coming up in Windows using hard drive.

Hard Drive Failure Symptoms & Associated Hard Drive Parts			
a)	Noise	PCB Board Component Failure Actuator Assembly/Head Failure DC Drive Spindle Motor Failure	} Logical parts, Physical parts failure
b)	Heat	DC Drive Spindle Motor Failure Spindle Bearing Wear Hard Drive Platter Misalignment	} Mostly physical failures
c)	Smoke	PCB Board Component Failure Windows Corruption Data Fragmentation	} Mostly logical but some physical failures
d)	Access	PCB Board Component Failure Windows Corruption Actuator Assembly/Head Failure DC Drive Spindle Motor Failure	} Logical parts, Physical parts failure
e)	Errors	Windows Corruption Data Fragmentation PCB Board Component Failure	} Mostly logical but some physical failures

VIII. HARD DISK CAPACITY & SYSTEM USES

The capacity of an HDD reported by measuring calculation and redundancy of the capacity:

A. Calculation

The gross capacity of older Hard disk drives can be calculated by multiplying for each zone of the drive the number of cylinders by the number of heads by the number of sectors/zone by the number of bytes/sector (most commonly 512) and then summing the totals for all zones. Some modern SATA drives will also report cylinder-head-sector (C/H/S) values to the CPU but they are no longer actual physical parameters since the reported numbers are constrained by historic operating-system interfaces.

The old C/H/S scheme has been replaced by logical block addressing. In some cases, to try to "force-fit" the C/H/S scheme to large-capacity drives, the number of heads was given as 64, although no modern drive has anywhere near 32 platters: the typical 2 TB hard disk as of 2013 has two 1 TB platters (and 4 TB drives use four platters).

B. Redundancy

In modern Hard disk drives spare capacity for defect management is not included in the published capacity; however in many early Hard disk drives a certain number of sectors were reserved for spares, thereby reducing capacity available to end users. In some systems, there may be hidden partitions used for system recovery that reduce the capacity available to the end user. For RAID subsystems, data integrity and fault-tolerance requirements also reduce

the realized capacity. For example, a RAID1 subsystem will be about half the total capacity as a result of data mirroring. RAID5 subsystems with x drives, would lose 1/x of capacity to parity. RAID subsystems are multiple drives that appear to be one drive or more drives to the user, but provides a great deal of fault-tolerance. Most RAID vendors use some form of checksums to improve data integrity at the block level. For many vendors, this involves using Hard disk drives with sectors of 520 bytes per sector to contain 512 bytes of user data and eight check sum bytes or using separate 512-byte sectors for the check sum data.

C. SYSTEM USE

The presentation of an HDD to its host is determined by its controller. The actual presentation may differ substantially from the drive's native interface, particularly in mainframes servers. Modern Hard disk drives, such as SAS and SATA drives, appear at their interfaces as a contiguous set of logical blocks that are typically 512 bytes long, though the industry is in the process of changing to the 4,096-byte logical blocks layout, known as the Advanced Format (AF).

The process of initializing these logical blocks on the physical disk platters is called low-level formatting, which is usually performed at the factory and is not normally changed in the field. As a next step in preparing an HDD for use, high-level formatting writes partition and file system structures into selected logical blocks to make the remaining logical blocks available to the host's operating system and its

applications. The file system uses some of the disk space to structure the HDD and organize files, recording their file names and the sequence of disk areas that represent the file. Examples of data structures stored on disk to retrieve files include the File Allocation Table (FAT) in the DOS file

system and inodes in many UNIX file systems, as well as other operating system data structures (also known as metadata). As a consequence, not all the space on an HDD is available for user files, but this system overhead is usually negligible.

Unit prefixes					
Advertised capacity by manufacturer (using decimal multiples)		Expected capacity by consumers in class action (using binary multiples)		Reported capacity	
With prefix	Bytes	Bytes	Diff.	Windows (using binary multiples)	Mac OS X 10.6+ (using decimal multiples)
100 <u>MB</u>	100,000,000	104,857,600	4.86%	95.4 MB	100 MB
100 <u>GB</u>	100,000,000,000	107,374,182,400	7.37%	93.1 GB, 95,367 MB	100 GB
1 <u>TB</u>	1,000,000,000,000	1,099,511,627,776	9.95%	931 GB, 953,674 MB	1,000 GB, 1,000,000 MB

So, the total capacity of Hard disk drives is given by manufacturers in megabytes (1 MB = 1,000,000 bytes), gigabytes (1 GB = 1,000,000,000 bytes) or terabytes (1 TB = 1,000,000,000,000 bytes). This numbering convention, where prefixes like "mega" and "Giga" denote powers of 1,000, is also used for data transmission rates and DVD capacities. However, the convention is different from that used by manufacturers of memory (RAM, ROM) and Compact disk drives, where prefixes like "kilo" and "mega" mean powers of 1,024. The practice of using prefixes assigned to powers of 1,000 within the HDD and computer industries dates back to the early days of computing. Computers do not internally represent HDD or memory capacity in powers of 1,024; reporting it in this manner is just a convention. Microsoft Windows uses the powers of 1,024 convention when reporting HDD capacity, thus an HDD offered by its manufacturer as a 1 TB drive is reported by these cases as a 931 GB HDD. Mac OS X 10.6 ("Snow Leopard"), uses powers of 1,000 when reporting HDD capacity.

Defragmentation is a procedure used to minimize delay in retrieving data by moving related items to physically proximate areas on the disk. Some computer operating systems perform defragmentation automatically. Although automatic defragmentation is intended to reduce access delays, performance will be temporarily reduced while the procedure is in progress.

a. *Seek time*

An Hard disk drives *Average Access Time* is its average seek time which technically is the time to do all possible seeks divided by the number of all possible seeks, but in practice is determined by statistical methods or simply approximated as the time of a seek over one-third of the number of tracks. Average seek time ranges from under 4 ms for high-end server drive¹ to 15 ms for mobile drives, with the most common mobile drives at about 12 ms and the most common desktop type typically being around 9 ms. The first HDD had an average seek time of about 600 ms; by the middle of 1970s Hard disk drives were available with seek times of about 25 ms. Some early PC drives used a stepper motor to move the heads, and as a result had seek times as slow as 80–120 ms, but this was quickly improved by voice coil type actuation in the 1980s, reducing seek times to around 20 ms. Seek time has continued to improve slowly over time. Faster seek rates typically require more energy usage to quickly move the heads across the platter, causing louder noises from the pivot bearing and greater device vibrations as the heads are rapidly accelerated during the start of the seek motion and decelerated at the end of the seek motion. Quiet operation reduces movement speed and acceleration rates, but at a cost of reduced seek performance.

b. *Rotational Latency*

Time to access data can be improved by increasing rotational speed (thus reducing latency) and/or by reducing the time spent seeking. Increasing areal density increases throughout by increasing data rate and by increasing the amount of data under a set of heads, thereby potentially reducing seek activity for a given amount of data. The time to access data has not kept up with throughput increases, which themselves have not kept up with growth in bit density and storage capacity.

IX. PERFORMANCE METRICES

To judge the performance of the working of the hard disk drive the following considerations will affect:

- a. Seek Time
- b. Rotational Latency
- c. Data Transfer Rate

The factors that limit the time to access the data on an HDD are mostly related to the mechanical nature of the rotating disks and moving heads. Seek time is a measure of how long it takes the head assembly to travel to the track of the disk that contains data. Rotational latency is incurred because the desired disk sector may not be directly under the head when data transfer is requested. These two delays are on the order of milliseconds each. The bit rate or data transfer rate (once the head is in the right position) creates delay which is a function of the number of blocks transferred; typically relatively small, but can be quite long with the transfer of large contiguous files. Delay may also occur if the drive disks are stopped to save energy.

Rotational speed [rpm]	Average latency [ms]
15,000	2
10,000	3
7,200	4.16
5,400	5.55
4,800	6.25

Latency is the delay for the rotation of the disk to bring the required disk sector under the read-write mechanism. It depends on rotational speed of a disk, measured in revolutions per minute (rpm). Average rotational latency is shown in the table on the right, based on the statistical relation that the average latency in milliseconds for such a drive is one-half the rotational period.

c. Data transfer rate

As of 2010, a typical 7,200-rpm desktop HDD has a sustained "disk-to-buffer" data transfer rate up to 1,030 Mbits/sec. This rate depends on the track location; the rate is higher for data on the outer tracks (where there are more data sectors per rotation) and lower toward the inner tracks (where there are fewer data sectors per rotation); and is generally somewhat higher for 10,000-rpm drives. A current widely used standard for the "buffer-to-computer" interface is 3.0 Gbit/s SATA, which can send about 300 megabyte/s (10-bit encoding) from the buffer to the computer, and thus is still comfortably ahead of today's disk-to-buffer transfer rates. Data transfer rate (read/write) can be measured by writing a large file to disk using special file generator tools, then reading back the file. Transfer rate can be influenced by file system fragmentation and the layout of the files.

HDD data transfer rate depends upon the rotational speed of the platters and the data recording density. Because heat and vibration limit rotational speed, advancing density becomes the main method to improve sequential transfer rates. Higher speeds require a more powerful spindle motor, which creates more heat. While areal density advances by increasing both the number of tracks across the disk and the number of sectors per track, only the latter increases the data transfer rate for a given rpm. Since data transfer rate performance only tracks one of the two components of areal density, its performance improves at a lower rate.

X. MARKET SEGMENTS

a. HARD DISK DRIVES (Mobile laptop)

They are smaller than their desktop and 60 GB and 4 TB and rotate at 5,400 to 10,000 rpm, and have a media transfer rate of 0.5 Gbit/s or higher (1 GB = 10^9 bytes; 1 Gbit/s = 10^9 bit/s). As of August 2014, the highest-capacity desktop Hard disk drives store 8 TB.

b. Desktop HARD DISK DRIVES

They typically store between enterprise counterparts, tend to be slower and have lower capacity. Mobile Hard disk drives spin at 4,200 rpm, 5,200 rpm, 5,400 rpm, or 7,200 rpm, with 5,400 rpm being typical. 7,200 rpm drives tend to be more expensive and have smaller capacities, while 4,200 rpm models usually have very high storage capacities. Because of smaller platter(s), mobile Hard disk drives generally have lower capacity than their greater desktop

c. Enterprise HARD DISK DRIVES

Typically used with multiple-user computers running enterprise software. Examples are: transaction processing databases, internet infrastructure (email, web server, e-commerce), scientific computing software, and near line storage management software. Enterprise drives commonly operate continuously ("24/7") in demanding environments while delivering the highest possible performance without sacrificing reliability. Maximum capacity is not the primary goal, and as a result the drives are often offered in capacities that are relatively low in relation to their cost. The fastest enterprise Hard disk drives spin at 10,000 or 15,000 rpm, and can achieve sequential media transfer speeds above 1.6 Gbit/s¹ and a sustained transfer rate up to 1 Gbit/s. Drives running at 10,000 or 15,000 rpm use smaller platters to mitigate increased power requirements (as they have less air drag) and therefore generally have lower capacity than the highest capacity desktop drives. Enterprise Hard disk drives are commonly connected through Serial Attached SCSI (SAS) or Fibre Channel (FC). Some support multiple ports, so they can be connected to a redundant host bus adapter. They can be reformatted with sector sizes larger than 512 bytes (often 520, 524, 528 or 536 bytes). The additional storage can be used by hardware RAID cards or to store a Data Integrity Field.

d. Consumer electronics HARD DISK DRIVES

They include drives embedded into digital video recorders and automotive vehicles. The former are configured to provide a guaranteed streaming capacity, even in the face of read and write errors, while the latter are built to resist larger amounts of shock.

e. External hard drives



Figure 14: Toshiba 1 TB 2.5" external USB 2.0 HDD

External HDDs typically connect via USB; variants using USB 2.0 interface generally have slower data transfer rates when compared to internally mounted hard drives connected through SATA. Plug and play drive functionality offers system compatibility and features large storage options and portable design. External HDDs are usually available as pre-assembled integrated products, but may be also assembled by combining an external enclosure (with USB or other interface) with a separately purchased HDD. They are available in 2.5-inch and 3.5-inch sizes; 2.5-inch variants are typically called *portable external drives*, while 3.5-inch variants are referred to as *desktop external drives*. "Portable" drives are packaged in smaller and lighter

enclosures than the "desktop" drives; additionally, "portable" drives use power provided by the USB connection, while "desktop" drives require external power bricks.

As from April 2014, capacities of external Hard disk drives generally range from 160 GB to 6 TB; common sizes are 160 GB, 250 GB, 320 GB, 500 GB, 640 GB, 750 GB, 1 TB, 2 TB, 3 TB, 4 TB, 5 TB and 6 TB. Features such as biometric security or multiple interfaces (for example, Firewire) are available at a higher cost. There are pre-assembled external hard disk drives that, when taken out from their enclosures, cannot be used internally in a laptop or desktop computer due to embedded USB interface on their printed circuit boards, and lack of SATA (or Parallel ATA) interfaces.

XI. FUTURE DEVELOPMENT

Modern drives make extensive use of error correction codes (ECCs), particularly Reed-Solomon error correction. These techniques store extra bits, determined by mathematical formulas, for each block of data; the extra bits allow many errors to be corrected invisibly. The extra bits themselves take up space on the HDD, but allow higher recording densities to be employed without causing uncorrectable errors, resulting in much larger storage capacity. For example, a typical 1 TB hard disk with 512-byte sectors provides additional capacity of about 93 GB for the ECC data. DC Drive Spindle Motor Failure (part Logical, part Physical failures).

The S.M.A.R.T (Self-Monitoring, Analysis and Reporting Technology) feature counts the total number of errors in the entire HDD fixed by ECC (although not on all hard drives as the related S.M.A.R.T attributes "Hardware ECC Recovered" and "Soft ECC Correction" are not consistently supported), and the total number of performed sector remapping, as the occurrence of many such errors may predict an HDD failure.

The worst type of errors are those that go unnoticed, and are not even detected by the disk firmware or the host operating system. These errors are known as silent data

A. NEW PARTICLE MATERIAL FOR 20 TIMES DENSER HARD DRIVES





Figure 15: SKYRMIONS

A groundbreaking new particle has been discovered that could potentially be used in extremely dense hard drives, known as SKYRMIONS or microscopic twisted magnetic vortices, the particle could prove extremely useful in the future of storage on magnetic hard drives, where traditional drives fast approach a density barrier.

B. HISTORY OF SKYRMIONS

corruption. Increasing areal density corresponds to an ever decreasing bit cell size. In 2013 a production desktop 3 TB Byte HDD (4 platters) would have had an areal density of about 500 G bit/in² which would have amounted to a bit cell comprising about 18 magnetic grains (11 by 1.6 grains). Since the mid-2000s areal density progress has increasingly been challenged by a super paramagnetic trilemma involving grain size, grain magnetic strength and ability of the head to write. In order to maintain acceptable signal to noise smaller grains are required; smaller grains may self-reverse (thermal instability) unless their magnetic strength is increased, but known write head materials are unable to generate a magnetic field sufficient to write the medium. Several new magnetic storage technologies are being developed to overcome or at least abate this trilemma and thereby maintain the competitiveness of Hard disk drives with respect to products such as flash memory-based solid-state drives (Secondary storage devices).

One such technology, shingled magnetic recording (SMR), was introduced in 2013 by Seagate as "the first step to reaching a 20 TB HDD by 2020"; however, SMR comes with design complexities that may cause slower write performance. Other new recording technologies that, as of 2014, still remain under development include heat-assisted magnetic recording (HAMR), microwave-assisted magnetic recording (MAMR), two-dimensional magnetic recording (TDMR), bit-patterned recording (BPR), and "current perpendicular to plane" giant magneto resistance (CPP/GMR) heads.

Depending upon assumptions on feasibility and timing of these technologies, the median forecast by industry observers and analysts for 2016 and beyond for areal density growth is 20% per year with a range of 10% to 40%. The ultimate limit for the BPR technology may be the super paramagnetic limit of a single particle that is estimated to be about two orders of magnitude higher than the 500 G bits/in² density represented by 2013 production desktop Hard disk drives .

Skyrmions are discovered after the UK particle physicist Tony Skyrme, who in 1962 found that they could explain how sub atomic particles such as neutrons and protons exist as discrete entities emerging from a continuous nuclear field. Scientists have known about skyrmions since the 1960s, but a recent discovery has allowed them to create and destroy them at will. Skyrmions occur in many materials including manganese-silicide thin films (in which they were discovered) and cobalt-iron-silicon.



Figure 16: Skyrmions

C. CREATING AND ANNIHILATING SINGLE SKRMIONS

In the new work the researchers studied a palladium-iron bilayer on an iridium-crystal surface. The tiny vertices can be imagined as 2D knots in which the magnetic moment rotates about 360° within a plane. The technique involves using a scanning tunnelling microscope (STM) and a polarized current to force group of atoms into knot-like twisted configurations. The skyrmions resist unravelling, and so they are theoretically used to store zeros (untwisted) or ones (twisted), the key to storing data digitally. To write or delete a skyrmion we position our STM tip, at a particular spot on the sample and inject a spin polarized tunnel-current pulse into it. While at low currents and voltages the sample magnetization is stable, at higher currents and voltages the magnetic state starts to switch between a skyrmion and a simple parallel alignment of magnetic moments. In this situation, the current direction is favoured over the other.

By using **Skyrmionic** hard disks could store 20 times more data per unit area than current hard disks, as each skyrmion is only a few nanometers in diameter. This means that instead of having a hard drive with 4 TB of space, in the future a skyrmionic drive of same size could hold 80 TB. So, this would lead to a drastic improvement in the storing capacity so for the future perceptible it gives a much better performance in the technology of the Hard Disk Drive.

XII. APPLICATIONS OF HARD DISK

A. SCAREWARE DETECTED

A fresh kind of scare ware, which pretends to be worthwhile application for Hard –disks, is circulating online similar to a bogus anti-spyware scam. Bogus anti-spyware attempts at making user believe that spyware has infected their PCs, phony hard –disk benefiting application attempts at making users believe that their PCs are disintegrating into parts.

B. 3D OPTICAL DATA STORAGE

3D optical data storage is the term given to any form of optical data storage in which information can be recorded and/or read with three-dimensional resolution (as opposed to the two-dimensional resolution afforded, for example, by CD). This innovation has the potential to provide petabyte-level mass storage on DVD-sized discs (120mm). Data recording and read back are achieved by focusing lasers within the medium. However, because of the volumetric nature of the data structure, the laser light must travel through other data points before it reaches the point where reading or recording is desired. Therefore, some kind of nonlinearity is required to ensure that these other data points do not interfere with the addressing of the desired point. No commercial product based on 3D optical data storage has yet arrived on the mass market, although several companies are actively developing the technology and claim that it may become available "soon". Layer-




Figure 17: Schematic representation of a cross-section through a 3D optical storage disc (yellow) along a data track (orange marks).

By this figure, four data layers are seen, with the laser currently addressing the third from the top. The laser passes through the first two layers and only interacts with the third, since here the light is at a high intensity.

C. APPLICATIONS OTHER THAN COMPUTER

The amount of data requiring digital storage continues to grow at a fantastic rate. No longer confined to use in computers, hard disk drives can now be found in televisions, games consoles and other home entertainment systems as a store of non-volatile data. The amount of platinum and ruthenium has increased rapidly over the years as storage densities have increased to meet demand. It had discovered that wide range of ferromagnetic materials the direction of magnetization can be completely controlled by polarized light without the need for magnetic fields, a finding that could significantly affect the data memory and storage industries that produce hard disks and magnetic random access memories. It is focused on materials currently being developed for high-density storage applications.

Ferromagnetism's most familiar form is the humble refrigerator magnet, but it is also a core component in many electrical devices, including magnetic storage used in commercial computing applications. In traditional magnetic storage devices magnetic bits are switched using magnetic fields, a slow process that consumes considerable energy and is reaching its density limits. This results showing that it is possible to switch magnetic bits using only the polarization of light could significantly simplify the design and improve the speed of magnetic recording. Magnetic storage is emerging in the memory market due to demands for higher-density, fast, and low power non-volatile memory. As industry trends toward silicon nano photonics, miniaturization, and photonic-electronic integration the ability to couple photonic, electronic, and magnetic materials could enable completely new applications.

It had been tested a rapid-pulse laser at a variety of ferromagnetic materials including magnetic thin films, multilayers and granular films. Previously, scientists have only been able to use all-optical control on a small set of ferrimagnetic materials that did not lend themselves for data storage applications. The next step is to scale the technology

to be able to write data on the nanoscale and time scales required for magnetic recording.

XIII. CONCLUSION

So, far we had discussed about the current status and future perspective of magnetic behavior analysis for the hard disk drive. The hard disk drive has proved as pivotal as the printing press in terms of its impact on human society. This paper examines both the history of and the prospects for the technology supporting these remarkable devices. Certainly for the next few years, conventional magnetic recording will continue its advance and provide ever higher capacities. There are several new technology features such as perpendicular recording and secondary actuators that are only now being introduced and which will drive further increases in a real density as they are refined and brought to maturity. However conventional magnetic recording appears limited to about 1 Tbit/in². Beyond this density, exciting new approaches such as heat-assisted magnetic recording or patterned magnetic media will be brought to bear.


With the increase in the recording density of Hard disk drives, the narrower magnetic pole and read sensor widths in writing and reading heads are required, and techniques used for the magnetic behaviour analysis in a nanoscale area is essential. The observation of the in-plane domain structure of the writer pole for a perpendicular recording (PMR) head was performed by electron holography in order to elucidate

the mechanism of the pole erasure originating from the instability of the magnetic domain state.

The results revealed that the stability of the domain structure is strongly related to the domain wall trapping, and the pole erasure can be suppressed by realizing the stable domain structure. With respect to a current perpendicular to plane (CPP)-giant magneto-resistive (GMR) head, which has been providing a reading head element, the insertion of non-magnetic materials, such as Cu, between the ferromagnetic pinned and free layers is known to increase the MR ratio. For realizing the practical use of CPP-GMR head, effects of the insertion of materials on the increase in the MR ratio were investigated by X-ray magnetic circular dichroism (XMCD). The results revealed that XMCD is a powerful technique to obtain information on electronic states and magnetic moments with the help of the theoretical electronic band calculation.

Magnets get flipped by light

A quick laser zap can alter the magnetism of a host of materials. The finding opens up the possibility of using lasers to speed up and simplify data storage in computers. Computer hard drives read and write data by flipping the magnetization of memory cells with magnetic fields. It had discovered that for a limited set of materials, they could replace the relatively slow, energy-hogging magnetic fields with laser light. The clockwise or counter clockwise polarization of the light determined the direction of the materials' magnetization.



Images of magnetic domains in a cobalt platinum (Co/Pt) alloy multilayer film exposed to laser light where dark gray indicates one magnetization orientation, while the light gray indicates an opposite orientation. The images show that the final direction of magnetization can be controlled using circularly polarized light without the use of magnetic fields.

Figure 17: Schematic representation of a magnetic domains in a Cobalt Platinum

One material, a film containing iron and platinum, is already being tested in hard drive prototypes that use lasers as a heat source to make storing data easier. Fullerton's research suggests that upgrading the lasers to emit quick polarized pulses could enable the laser light to read and write data as well.

REFERENCES

- [1] "IBM Archives: IBM 350 disk storage unit". Retrieved 2012-10-19.
- [2] Fullerton, Eric (March 2014). "5th Non-Volatile Memories Workshop (NVMW 2014)". IEEE- Retrieved April 23, 2014.
- [3] "Time Capsule, 1956 Hard Disk". Oracle Magazine. Oracle. July 2014. Retrieved September 19, 2014. IBM 350 disk drive held 3.75 MB
- [4] "Toshiba Storage Solutions – MK3233GSG".
- [5] "Magnetic head development". IBM Archives. Retrieved August 11, 2014.
- [6] "Seagate ST5000DM000 Desktop HDD Product Manual". Seagate. Retrieved August 11, 2014.
- [7] Computer Organization and Design, 2nd Ed., P. Pal Chaudhuri.
- [8] New Paradigms in Magnetic Recording.